



AP-831

**APPLICATION
NOTE**

**Pentium® II Xeon™
Processor SMBus
Thermal Reference
Guidelines**

June 1998

Order Number: 243791-001

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Pentium® II Xeon™ processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an ordering number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com>

Copyright © Intel Corporation 1998.

* Third-party brands and names are the property of their respective owners.

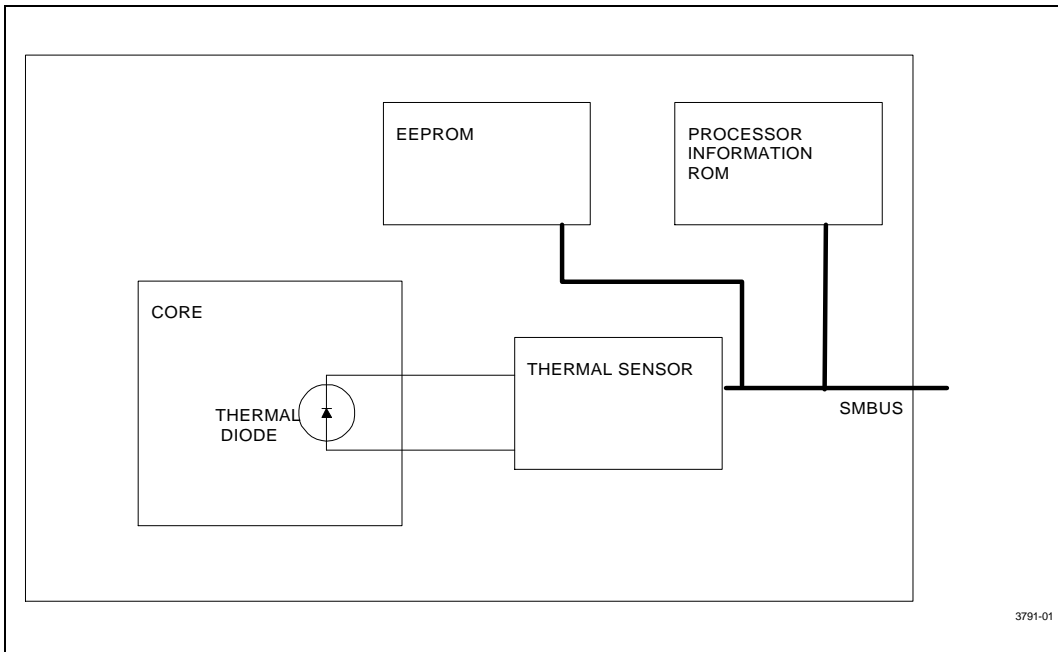
CONTENTS

	PAGE		PAGE
1.0. INTRODUCTION	4	3.3.3. SYSTEM MANAGEMENT SOFTWARE.....	6
1.1. Definition of Terms	5	3.3.4. SMBALERT# AS AN EARLY DETECTION MECHANISM.....	7
2.0. THE THERMAL SENSOR	5	3.4. Filtering the Thermal Sensor Readings.....	7
3.0. THERMAL REFERENCE BYTE	5	3.4.1. STATISTICAL ANALYSIS.....	7
3.1. Using the OEM Scratch EEPROM for Thermal Management	6	3.4.2. CONTROL THEORY.....	7
3.2. Calibration.....	6	3.4.3. RESPONSE TIME TO INITIATE SHUTDOWN.....	7
3.2.1. PROCESSOR LEVEL.....	6	4.0. THE PROCESSOR CORE THERMAL DIODE.....	7
3.2.2. SYSTEM LEVEL	6	4.1. Ideality Factor	7
3.3. Application Examples	6	4.2. Thermal Diode vs. Thermal Plate Measurements.....	8
3.3.1. STANDARD OPERATION.....	6	5.0. CONCLUSION	9
3.3.2. CUSTOMIZED THRESHOLD BYTES FOR SYSTEM COOLING DEPENDENCE.....	6		

1.0. INTRODUCTION

The Pentium® II Xeon™ processor, like the Pentium II processor, is housed in a Single Edge Contact (S.E.C.) cartridge. Intel has included a thermal sensor inside the Pentium II Xeon processor to aid in thermal management of server and workstation systems. This thermal sensor is connected to a thermal diode on the processor core. The core is the most significant source of power dissipation in the cartridge. This feature allows for flexibility and precision in thermal management applications. The thermal diode temperature changes with the junction temperature very closely. Thus, the thermal sensor provides the earliest indication of thermal variation and the onset of potential thermal solution failure. *Note that the maximum thermal plate specification remains the governing condition for guaranteed operation of the processor.* For some applications the cache may be the forcing function of the thermal plate temperature if the core is not being exercised in high power mode. The thermal diode is located on the processor core and does not monitor the cache temperature.

The basic use of the thermal sensor is described in the *Pentium® II Xeon™ Processor at 400 MHz* datasheet. The thermal sensor is provided to deliver thermal data to a system with which system management software can forecast the onset of a potentially catastrophic event and gracefully shutdown or take other reactive measures in order to preserve data integrity. No calibration is required to effectively apply this technology. If desired, custom threshold values may be determined for individual operating conditions. Example applications are discussed in this document. Additionally, statistical methods are discussed for effective filtering of thermal byte readings. Filtering ensures that isolated harmless transient changes in temperature are not construed as catastrophic conditions. This document describes the thermal reference feature in more detail, and includes suggestions and procedures for more effective use of the Pentium II Xeon processor Thermal Reference Byte and the thermal sensor in thermal management applications. The below figure is a schematic representation of the components relative to the thermal reference feature.



1.1. Definition of Terms

- *EEPROM* - electronically erasable, programmable read-only memory; a memory device that can be used by the OEM to store additional information, including information for a thermal solution.
- *Processor Information ROM* - a memory device where the Thermal Reference Byte and other information is stored.
- *SMBus* - System Management bus.
- *Thermal byte reading* - a reading from the thermal sensor. This byte represents core thermal data at the instant the reading is taken.
- *Threshold byte value* - a threshold value customized to particular operating conditions.
- *Thermal diode* - a diode located on the processor core which enables thermal measurements.
- *Thermal Reference Byte* - a byte value that corresponds to the maximum operating T_{PLATE} at maximum steady state power specifications for which the product is guaranteed to operate. This byte is determined during manufacturing test for each processor unit and is stored in the Processor Information ROM.
- *Thermal sensor* - a device which performs analog to digital conversions and supplies a single byte reading of thermal data (from the thermal diode).

2.0. THE THERMAL SENSOR

The Pentium II Xeon processor thermal sensor provides a means of acquiring thermal data from the processor core with an exceptional degree of precision. The thermal sensor is composed of control logic, SMBus interface logic, a precision analog-to-digital converter, and a precision current source. The thermal sensor drives a small current through the p-n junction of a thermal diode located on the same silicon die as the processor core. The forward bias voltage generated across the thermal diode is sensed and the precision A/D converter derives a single byte of data, or “thermal byte reading.” A/D conversions happen continuously and the single byte result of the most recent conversion is stored in a register on the thermal sensor. This single byte result can be read from the register using the serial interface protocol of the SMBus. A thermal byte reading for the thermal diode is obtained using an SMBus read command. The rate of reads is limited by the clock and the protocol of the SMBus serial interfacing. System management software

running on the processor or on a microcontroller can acquire the data from the thermal sensor for system management.

Upper and lower thermal reference thresholds can be individually programmed for the thermal diode. The register where the single byte result is stored is also sampled by comparator circuits. These circuits compare the single byte result against programmable threshold bytes. The alert signal on the Pentium II Xeon processor SMBus (SMBALERT#) will assert when either threshold is crossed.

The thermal byte reading can also be used in conjunction with the Thermal Reference Byte in the Processor Information ROM. Byte 9 of the Processor Information ROM contains the address in the ROM of this byte, described in more detail in Section 3.0. The thermal byte reading from the thermal sensor can be compared to this Thermal Reference Byte to provide an indication of the difference between the temperature of the processor core at the instant of the thermal byte reading and the temperature of the processor core under the steady state conditions of high power and maximum T_{PLATE} specifications. The nominal precision of the least significant bit of a thermal byte is 1 °C.

3.0. THERMAL REFERENCE BYTE

To increase the usefulness of the thermal diode and thermal sensor, Intel has added a new procedure to the manufacturing and test flow of the Pentium II Xeon processor. This procedure determines the Thermal Reference Byte and programs it into the Processor Information ROM. The Thermal Reference Byte is uniquely determined for each unit. The procedure causes each unit to dissipate its maximum power (which can vary from unit to unit) while at the same time maintaining the thermal plate at its maximum specified operating temperature. Correctly used, this feature permits an efficient thermal solution while preserving data integrity.

The value of the Thermal Reference Byte is determined individually for each unit during manufacturing test. While the processor runs code designed to draw high power and the thermal plate is forced to its maximum specified temperature, readings are taken from the thermal sensor to determine the Thermal Reference Byte. The expected value of this byte in the Processor Information ROM falls within a wide range. Contributing factors to this wide range of values include the variability between processors of power consumption, ideality factors (see Section 4.1.), and junction to plate thermal resistance (θ_{JP}).

3.1. Using the OEM Scratch EEPROM for Thermal Management

The Thermal Reference Byte provided in the Processor Information ROM by Intel can be used to calculate additional information for a thermal solution, which can then be stored in the OEM Scratch EEPROM, if desired. The Thermal Reference Byte need not be used as it exists in the Processor Information ROM. Other creative applications of this feature may include deriving an alternative (threshold byte value) which is better tuned to a particular application. A discussion of these applications is provided in Section 3.3.

3.2. Calibration

There is no calibration required to effectively apply this technology. Traditional methods of calibration against an absolute temperature scale may be applied with conventional methods. However, this is not necessary as the method used for programming the Thermal Reference Byte incorporates a built-in calibration for process variation.

3.2.1. PROCESSOR LEVEL

Dimensionless byte values corresponding to meaningful system application conditions provide a direct and efficient method of establishing thresholds for thermal management software use. Intel preprograms a dimensionless threshold unique to each processor in the form of the Thermal Reference Byte. Section 3.3. describes applications that use the Thermal Reference Byte and real time thermal byte readings, which can be applied to produce superior total thermal solutions.

3.2.2. SYSTEM LEVEL

Thermal byte readings at nominal and/or maximum system operating conditions may be taken and compared to the Thermal Reference Byte. This delta is proportional to the temperature difference. The customized information can be programmed into any nonvolatile memory, including the OEM Scratch EEPROM, and used as a basis for operations such as thermal sensor limit programming and shutdown action. Examples are described further in Section 3.3.

3.3. Application Examples

Note that in the application examples described in Sections 3.3.1. through 3.3.4., the maximum thermal plate temperature remains the governing condition for guaranteed correct functionality of any S.E.C. cartridge processor.

3.3.1. STANDARD OPERATION

It is simplest to take thermal byte readings and compare them to the Thermal Reference Byte. As the delta (Δ) between these readings approaches zero, actions can be taken to gracefully shut down or take other system-level action to preserve data integrity. The proportional delta, as well as the rate of change (derivative), or even the cumulative sum (integral) of this delta may be used for a more sophisticated closed loop system for optimal thermal management.

3.3.2. CUSTOMIZED THRESHOLD BYTES FOR SYSTEM COOLING DEPENDENCE

Consider changing the above scenario by using system level cooling of the thermal plate which maintains the thermal plate temperature significantly below the maximum T_{PLATE} specification. Clearly, the thermal byte readings will not approach the Thermal Reference Byte until sometime after the onset of a potentially catastrophic event, and may be too late to begin safe system shutdown or other reactive procedures. Under these conditions, and assuming the application calls for the earliest possible detection of the onset of a potentially catastrophic event, establishing a threshold byte value customized to the unique operating environment and conditions is advised. Just as in the application described in Section 3.3.1., as the delta approaches zero actions can be taken to perform safe shutdown or other appropriate system-level action. This will allow the earliest response to the onset of a possible problem even for the particular case of a superior cooling system.

3.3.3. SYSTEM MANAGEMENT SOFTWARE

Through use of the system management features, systems can be designed to react sooner than that which was possible with previous thermal alert systems. System management software can be used to optimize the timing for detection of the onset of a potential problem for specific operating conditions. This software may be run at any time to update threshold byte values so that the individualized threshold is set at the appropriate value for the current operating conditions.

3.3.4. SMBALERT# AS AN EARLY DETECTION MECHANISM

In the above applications SMBus bandwidth is necessary to take thermal byte readings. To avoid wasteful use of bandwidth, an asynchronous interrupt signal, SMBALERT#, may be used. To avoid the need to poll each processor continuously, an early warning can be programmed into the thermal reference thresholds. This will cause the SMBALERT# signal to be sensed by the system management microcontroller or processor, and the early warning can be used to cause the system management hardware and software to begin monitoring thermal byte values. After a brief period of monitoring thermal byte readings, system management software can select whether the event that triggered the early warning alert was a potential problem event or a false alarm, and take corrective action or clear the alert as appropriate.

3.4. Filtering the Thermal Sensor Readings

3.4.1. STATISTICAL ANALYSIS

When using the thermal sensor to provide information to thermal management software, care should be taken to ensure that harmless transient changes in temperature are not construed as indicating the onset of catastrophic conditions. The quick response of the thermal diode to temperature changes in the processor core will inevitably yield isolated transient readings that could be misleading. It should not be presumed that an isolated byte reading is indicative of the onset of a catastrophic thermal failure. Performing data analyses or filtering is an effective tactic to avoid false alarms. Relevant statistics such as mean, standard deviation, mode, median, running average, minimum, maximum, etc., may be used. The choice of which combination of these statistics is optional and may be different for each system and application. Statistical averaging using mean and standard deviation is one effective method of filtering the thermal byte readings.

3.4.2. CONTROL THEORY

Different uses of the delta function may be of primary concern depending on the particular situation. If the temperature is slowly approaching the threshold, the value of primary concern is the delta. However, if the temperature is rapidly approaching the threshold one would want to monitor the rate ($d\Delta/dt$) in order to react quickly in anticipation of a possible problem depending on the value of Δ . Furthermore, if the rate ($d\Delta/dt$) is

relatively unchanging but the temperature is near the threshold, the time integral may be taken if it is desired to push the operating envelope. Indeed, the choice of proportional integral and differential gain factors is once again a customized choice particular to the application environment and the processor power consumption in response to the application software.

3.4.3. RESPONSE TIME TO INITIATE SHUTDOWN

Server management software which processes the thermal byte reading should perform heavy damping or filtering of several hundred or perhaps thousands of data points (depending on the speed of polling) before initiating graceful shutdown or other system-level action. The appropriate delay before responding will depend not only on the time required for a graceful response sequence but also on the time constant of the external thermal management system designed to limit T_{PLATE} to within the specified operating range.

Using system and even application specific characterization data, it is feasible for adaptive values to be established for coefficients of the filtering function. This will modulate the delayed response, effectively customizing it for the specific requirements of the system or application.

4.0. THE PROCESSOR CORE THERMAL DIODE

4.1. Ideality Factor

The thermal diode is on the processor core of the Pentium II Xeon processor. The V-I relationship of a p-n junction diode can be described by the ideal diode equation:

$$I = I_s \left(\exp \frac{eV}{kT} \right) - 1$$

where I_s is the reverse saturation current and kT/e is a volt-equivalent of temperature. Maximum deviation of a particular diode from ideal behavior is represented by an *ideality factor*. Ideality factors vary from diode to diode. The range of ideality factors is one reason why absolute temperature accuracy is not specified for readings from the thermal diode.

4.2. Thermal Diode vs. Thermal Plate Measurements

An intimate internal thermal connection exists between the on-core thermal diode and the most significant source of thermal power. The thermal diode temperature varies with the processor core junction temperature much more closely than the thermal plate, and can give the earliest indication of the direction in which the T_{PLATE} will move. This allows detection of thermal variation and the onset of thermal solution failure much sooner than a thermocouple attached to the thermal plate, or even a thermal diode within the cartridge itself. The thermal mass of the core package, thermal grease, thermal plate, and any heatsink or heat exchange device all contribute to a slower response time for thermal solutions that do not use the on-core thermal diode.

A slower response may mask a perturbation, which could be a potentially catastrophic event. Being the earliest indicator of a possible problem, the on-core thermal diode used in temperature measurements provides a tremendous benefit in thermal management. Keep in mind that the maximum thermal plate specification remains the governing condition for guaranteed operation of the processor. For some applications the cache may be the component driving the thermal plate temperature if the core is not being exercised in high power mode. The thermal diode is located on the processor core and does not monitor the cache temperature.

A complex relationship exists between the core junction temperature and the plate temperature in steady state. This relationship is expressed by the equation:


$$T_{\text{JUNCTION}} = T_{\text{PLATE}} + \theta_{\text{JP}} * \text{Power}.$$

In the transient state, the relationship is much more complex and is typically expressed by the summation of a series of step response functions. Each step response can be expressed by the equation:

$$T_x(t) = T_1 + (T_2 - T_1) \left(1 - e^{\left(\frac{-t}{\tau} \right)} \right)$$

where T_1 is the initial temperature, T_2 is the final temperature, and τ a time constant which varies with each thermal solution.

The below figure illustrates the slower response time of the thermal plate to indicate temperature compared to the temperature of the core itself. This thermal data was taken by attaching a thermocouple to the thermal plate directly above the contact fingers with thermal tape. A second thermocouple was placed directly on top of the processor core package and embedded between the core heat slug and the cartridge thermal plate. The cartridge temperature was forced to the maximum specified condition. Various typical operations were performed (as shown in the below figure) while thermal data was collected via a thermal data logger.



3791-02

As the above figure illustrates, the T_{PLATE} response lags T_{CORE} and the swings in T_{CORE} are typically larger than the swings in T_{PLATE} .

5.0. CONCLUSION

Thermal management applications are enabled to become more flexible by including a thermal sensor located on the processor. The internal thermal connection of the sensor to the die is intimate. The thermal diode temperature changes with the junction temperature much more closely than was possible for traditional thermal sensors. Thus, the thermal sensor on the core provides the earliest indication of thermal variation and the onset of thermal solution failure. This enables server management software to invoke a timely response to the onset of a potentially catastrophic event and gracefully shutdown or take other measures in order to preserve data integrity. **Keep in mind that the maximum thermal plate specification remains the governing condition for guaranteed operation of the processor.**

A microcontroller may respond to a convenient asynchronous interrupt generator provided by the thermal alert signal (SMBALERT#). This may potentially involve system management software to begin sampling or filtering thermal byte readings before determining whether shutdown or other system-level action is necessary. Many methods of filtering a series of readings have been described as effective strategies for discerning between harmless transient spikes in temperature and bona fide indications of a thermal solution failure.

The Thermal Reference Byte is present as a convenience feature and is programmed by Intel for each unit during manufacturing test. No calibration is required to effectively apply this technology. Custom threshold byte values may be determined for individual operating conditions if desired, and examples of applications for which this is appropriate have been discussed. These unique threshold bytes allow for earliest detection of a possible problem under specific operating conditions. Creative applications of this set of features are possible and encouraged, especially as they cater to a system's individual thermal solutions or application conditions.



UNITED STATES, Intel Corporation
2200 Mission College Blvd., P.O. Box 58119, Santa Clara, CA 95052-8119
Tel: +1 408 765-8080

JAPAN, Intel Japan K.K.
5-6 Tokodai, Tsukuba-shi, Ibaraki-ken 300-26
Tel: + 81-29847-8522

FRANCE, Intel Corporation S.A.R.L.
1, Quai de Grenelle, 75015 Paris
Tel: +33 1-45717171

UNITED KINGDOM, Intel Corporation (U.K.) Ltd.
Pipers Way, Swindon, Wiltshire, England SN3 1RJ
Tel: +44 1-793-641440

GERMANY, Intel GmbH
Dornacher Strasse 1
85622 Feldkirchen/ Muenchen
Tel: +49 89/99143-0

HONG KONG, Intel Semiconductor Ltd.
32/F Two Pacific Place, 88 Queensway, Central
Tel: +852 2844-4555

CANADA, Intel Semiconductor of Canada, Ltd.
190 Attwell Drive, Suite 500
Rexdale, Ontario M9W 6H8
Tel: +416 675-2438